workloads at times of the day or week during which content provisioning activities are scheduled.

A variable resource parameter may optionally be implemented to vary dynamically according to monitored information management system I/O resources/characteristics, such as monitored background system processing activity. For example, in one exemplary embodiment, processing background system activity may be monitored (*e.g.*, by monitoring arrival queue of background I/O requests to determine if the existing value of Reserved_Factor needs to be changed. Background system I/O activity includes, for example, write or update requests for new content, access to file system D-node/I-node data, and/or access to overhead blocks in a continuous or streaming file. If the background I/O queue increases in size, the value of Reserved_Factor may be increased, proportionally or using some other desired relationship. If the background I/O queue decreases in size, the value of Reserved_Factor may be decreased proportionally or using some other desired relationship. If the background I/O queue is empty, the value of Reserved_Factor may be set to zero. If desired, upper and/or lower bounds for Reserved_Factor (*e.g.* upper bound of about 0.05; lower bound of about 0.4) may be selected to limit the range in which the Reserved_Factor may be dynamically changed.

In one exemplary embodiment, Reserved_Factor may be dynamically varied from a first value ("Old_Reserved_Factor") to a second value ("New_Reserved_Factor") in a manner directly proportional to the background system activity workload. For example, Reserved_Factor may be dynamically varied from a first value ("Old_Reserved_Factor") to a second value ("New_Reserved_Factor") in a manner directly proportional to a change from a first monitored background system I/O queue size ("Old_Queue_Depth") to a second monitored background system I/O queue size ("New_Queue_Depth") using a proportionality factor ("C") and solving for the value "New_Queue_Depth" in the following equation:

$$\text{New\_Reserved\_Factor} - \text{Old\_Reserved\_Factor} = [C * (\text{New\_Queue\_Depth} - \text{Old\_Queue\_Depth})] \quad \text{(8B)}$$

In this exemplary embodiment, Old_Reserved_Factor may be a pre-determined initial value of Reserved_Factor set by, for example, operator or system manager input, or alternatively may be a value previously determined using equation (8B) or any other suitable equation or relationship.

It will be understood with benefit of this disclosure that equation (8B) is just one exemplary equation that may be employed to dynamically vary a resource parameter, such as Reserved_Factor, in a directly proportional manner with changing processing background activity. It will also be understood that the variables and constant "C" employed in equation (8B) are exemplary as well. In this regard, other equations, algorithms and/or relationships may be employed to dynamically vary the value of a resource parameter, such as Reserved_Factor, based on changes in background processing activity. For example, a resource parameter may be dynamically varied in a non-directly proportional manner, using other equations, algorithms or relationships, *e.g.*, using proportional ("P"), integral ("I"), derivative ("D") relationships or combinations thereof, such as proportional-integral ("PI"), derivative-integral-derivative ("PID"), *etc.* Furthermore, it will be understood that background processing activity may be measured or otherwise considered using alternative or additional factors to background I/O queue size, for example, by counting pending background I/O requests, *etc.*

Under certain circumstances, ongoing processing requirements (*e.g.*, for ongoing video streams) may be sufficiently high so that a newly determined resource parameter value such as New_Reserved_Factor may not be implemented without reducing interruption. To address this scenario, a dynamically-changing resource parameter embodiment such as previously described may be optionally implemented in a manner that controls rapid changes in parameter values to avoid interruptions to ongoing operations (*e.g.*, ongoing streams being viewed by existing viewers). For example, ongoing streams may be allowed to terminate normally prior to changing the value of Reserved_Factor, so that no interruption to existing viewers occur. This may be done, for example, by waiting until available processing resources are sufficient to increase the value of Reserved_Factor to New_Reserved_Factor, or by incrementally increasing the value of Reserved_Factor as streams terminate and additional processing resources become available.

Alternatively, processing background requirements may be given priority over service interruptions, in which case the existing streams of ongoing viewers may be immediately terminated as necessary to increase the Reserved_Factor to its newly determined value. In such an embodiment, existing streams may be selected for termination based on any desired factor or combination of such factors, such as duration of existing stream, type of viewer, type of stream, class of viewer, *etc.* In the latter case, lower classes of viewers may be terminated prior to higher class viewers and, if desired, some higher classes of viewers may be immune from termination as part of the guaranteed service terms of a service level agreement ("SLA"), other priority-indicative parameter (*e.g.,* CoS, QoS, *etc*), and/or other differentiated service implementation. Examples of possible SLA implementations, priority-indicative parameters and other differentiated service features may be found described in co-pending United States patent application serial number 09/879,810 filed on June 12, 2001 which is entitled SYSTEMS AND METHODS FOR PROVIDING DIFFERENTIATED SERVICE IN INFORMATION MANAGEMENT ENVIRONMENTS which is incorporated herein by reference

In operation, an I/O resource manager may utilize such a resource parameter to allocate cycle time T, for example, by using the parameter Reserved_Factor to determine a value of cycle T such that (1- Reserved_Factor) * T satisfies normal continuous file workload. Under such conditions, (1- Reserved_Factor) * T should be greater than or equal to the storage device service time, *i.e.,* the sum of access time and data transfer time. Accordingly, in this embodiment, cycle time T may be calculated to ensure sufficient I/O capacity for continuous playback for a number of viewers by using the following Resource Model Equations (9), (10) and (11) that correspond to respective Resource Model Equations (4), (7) and (8A).

For single storage device case:

$$NoV * AA / [1 - Reserved\_Factor - (\Sigma_{i=1}^{Nov} P_i) / TR] \leq T$$

$$\leq (1- Reserved\_Factor) * B_{max} / [(1 - B\_Save) * (\Sigma_{i=1}^{Nov} P_i)] \qquad (9)$$

For multiple storage device case under substantially balanced conditions: